

What I did to Adol-C and ISSM

Utke/Larour

Argonne National Laboratory
Jet Propulsion Laboratory

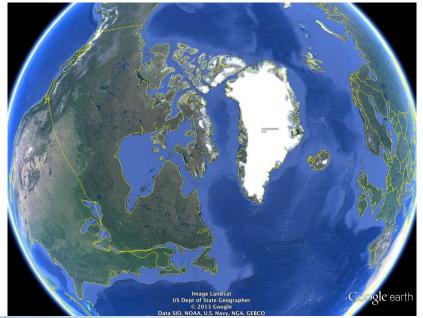
EuroAD - Dec/2013 Oxford, UK

outline

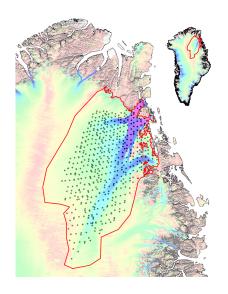
- what this is for
- changes in ISSM
- changes in Adol-C
- external solvers
- adjoinable MPI
- performance



Greenland



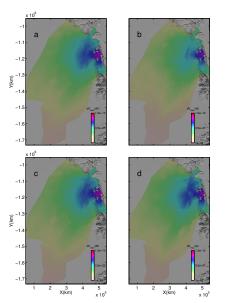
the North-Eastern Ice Stream on Greenland



- velocity field
- red boundary shows domain of interest
- dots indicate observation data
- surface observations by satellite/stations
- drilling holes is expensive
- goal is model tuning for prediction of sea level rise

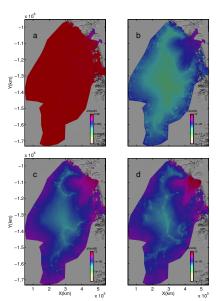


sensitivity studies - maximal velocity with respect to



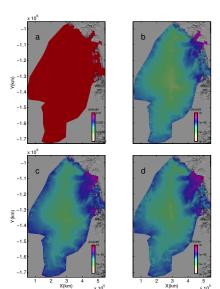
- a: *H* ice thickness
- \diamond b: S surface elevation
- \diamond c: B bed elevation
- \diamond d: α friction coefficient

sensitivity studies - volume with respect to



- a: *H* ice thickness
- \diamond b: S surface elevation
- \diamond c: B bed elevation
- \diamond d: α friction coefficient

sensitivity studies (last week) - volume with respect to

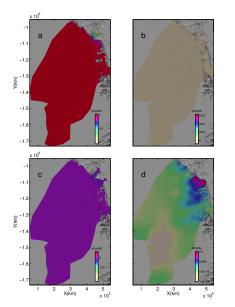


- a: *H* ice thickness
- \diamond b: S surface elevation
- ⋄ c: B bed elevation
- \diamond d: α friction coefficient

compared to earlier studies ran in higher resolution on Pleiades for longer model time

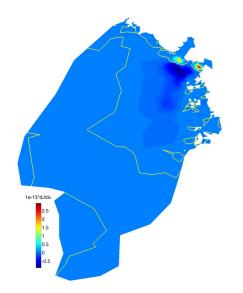


sensitivity studies - volume above floatation with respect to



- a: *H* ice thickness
- \diamond b: S surface elevation
- ⋄ c: B bed elevation
- \diamond d: α friction coefficient

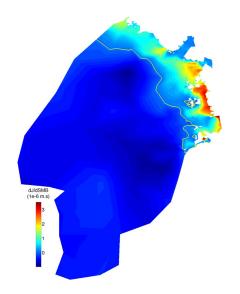
model-to-observation misfit of S to internal state



- with respect to friction coefficients
- \diamond L^2 integrated over time
- yellow lines indicate gradient sign switch
- ⋄ part of an gradient → line search optimization loop



model-to-observation misfit of S to external boundary



- with respect to snow-mass-balance
- snow fall given as (external) reanalysis of climate model runs
- hints at less snow fall on the coast, more inland
- means to adapt reanalysis if one assumes the ice sheet model is "correct"

presented at AGU meeting



◇ C++ model



- ◆ C++ model
- \diamond \approx 750 files, 100K lines, 10 developers



- ◆ C++ model
- $\diamond \approx 750$ files, 100K lines, 10 developers
- uses templates, virtual methods, multiple inheritance (i.e. isn't just glorified C)



- ◆ C++ model
- \diamond \approx 750 files, 100K lines, 10 developers
- uses templates, virtual methods, multiple inheritance (i.e. isn't just glorified C)
- only the model core is C++, data pre/post processing done with Matlab / Python



- ◆ C++ model
- \diamond \approx 750 files, 100K lines, 10 developers
- uses templates, virtual methods, multiple inheritance (i.e. isn't just glorified C)
- only the model core is C++, data pre/post processing done with Matlab / Python
- parallelized with MPI



- ◆ C++ model
- \diamond \approx 750 files, 100K lines, 10 developers
- uses templates, virtual methods, multiple inheritance (i.e. isn't just glorified C)
- only the model core is C++, data pre/post processing done with Matlab / Python
- parallelized with MPI



- ♦ C++ model
- \diamond \approx 750 files, 100K lines, 10 developers
- uses templates, virtual methods, multiple inheritance (i.e. isn't just glorified C)
- only the model core is C++, data pre/post processing done with Matlab / Python
- parallelized with MPI
- has extensive regression testing (incl. the numerical results)



- ♦ C++ model
- \diamond \approx 750 files, 100K lines, 10 developers
- uses templates, virtual methods, multiple inheritance (i.e. isn't just glorified C)
- only the model core is C++, data pre/post processing done with Matlab / Python
- parallelized with MPI
- has extensive regression testing (incl. the numerical results)
- uses libraries (meshing, partitioning, solvers)



typedef an IssmDouble and an IssmPDouble and switch on and off via _HAVE_ADOLC_ configure define



- typedef an IssmDouble and an IssmPDouble and switch on and off via _HAVE_ADOLC_ configure define
- contributors deliver code in terms of doubles and expert developer categorizes those into IssmDoubles and IssmPDoubles
 - \implies easy check



- typedef an IssmDouble and an IssmPDouble and switch on and off via _HAVE_ADOLC_ configure define
- contributors deliver code in terms of doubles and expert developer categorizes those into IssmDoubles and IssmPDoubles
 - ⇒ easy check
- replace all mallocs, news and frees, deletes by templated xNew / xDelete incl. variants for 2-D arrays => safer, cleaner, more efficient; easy check



- typedef an IssmDouble and an IssmPDouble and switch on and off via _HAVE_ADOLC_ configure define
- contributors deliver code in terms of doubles and expert developer categorizes those into IssmDoubles and IssmPDoubles
 - \implies easy check
- templatize data containers (partially done)
- some undue activation (still) forced through Matlab interface



- typedef an IssmDouble and an IssmPDouble and switch on and off via _HAVE_ADOLC_ configure define
- contributors deliver code in terms of doubles and expert developer categorizes those into IssmDoubles and IssmPDoubles
 - \implies easy check
- replace all mallocs, news and frees, deletes by templated xNew / xDelete incl. variants for 2-D arrays => safer, cleaner, more efficient; easy check
- templatize data containers (partially done)
- some undue activation (still) forced through Matlab interface
- passing data to passive code with templated reCast



- typedef an IssmDouble and an IssmPDouble and switch on and off via _HAVE_ADOLC_ configure define
- contributors deliver code in terms of doubles and expert developer categorizes those into IssmDoubles and IssmPDoubles
 - \implies easy check
- templatize data containers (partially done)
- some undue activation (still) forced through Matlab interface
- passing data to passive code with templated reCast
- reCast injections represent majority of the manual adaptation work



 pick a simple(!) setup to start with and establish consistency with FD tests



- pick a simple(!) setup to start with and establish consistency with FD tests
- few time steps, coarse resolution
- sequential, dense LU solve from GSL needs wrapping



- pick a simple(!) setup to start with and establish consistency with FD tests
- few time steps, coarse resolution
- sequential, dense LU solve from GSL needs wrapping
- Adol-C has had an "external" interface for func(x,y) all along but:



- pick a simple(!) setup to start with and establish consistency with FD tests
- few time steps, coarse resolution
- sequential, dense LU solve from GSL needs wrapping
- Adol-C has had an "external" interface for func(x,y) all along but:
 - the forward variants passed only $\dot{x},~\dot{y}$ and not x,~y themselves \Longrightarrow added



- pick a simple(!) setup to start with and establish consistency with FD tests
- few time steps, coarse resolution
- sequential, dense LU solve from GSL needs wrapping
- Adol-C has had an "external" interface for func(x,y) all along but:
 - the forward variants passed only \dot{x} , \dot{y} and not x, y themselves \Longrightarrow added
 - added sanity checks for consecutive locations



- pick a simple(!) setup to start with and establish consistency with FD tests
- few time steps, coarse resolution
- sequential, dense LU solve from GSL needs wrapping
- Adol-C has had an "external" interface for func(x,y) all along but:
 - the forward variants passed only \dot{x} , \dot{y} and not x, y themselves \Longrightarrow added
 - added sanity checks for consecutive locations
 - the forward/reverse handlers taped/restored x, y values whether needed or not
 - ⇒ added controls



- pick a simple(!) setup to start with and establish consistency with FD tests
- few time steps, coarse resolution
- sequential, dense LU solve from GSL needs wrapping
- Adol-C has had an "external" interface for func(x,y) all along but:
 - the forward variants passed only \dot{x} , \dot{y} and not x, y themselves \Longrightarrow added
 - added sanity checks for consecutive locations
 - the forward/reverse handlers taped/restored x, y values whether needed or not
 - ⇒ added controls
 - \bullet time stepping \to multiple solves with adaptive meshing \to changing system dimensions
 - ⇒ added tracking for maximum dimensions for single allocation of reusable help buffers



- pick a simple(!) setup to start with and establish consistency with FD tests
- few time steps, coarse resolution
- sequential, dense LU solve from GSL needs wrapping
- Adol-C has had an "external" interface for func(x,y) all along but:
 - the forward variants passed only \dot{x} , \dot{y} and not x, y themselves \Longrightarrow added
 - added sanity checks for consecutive locations
 - the forward/reverse handlers taped/restored x, y values whether needed or not
 - ⇒ added controls
 - time stepping → multiple solves with adaptive meshing → changing system dimensions
 - ⇒ added tracking for maximum dimensions for single allocation of reusable help buffers
- got first regression tests passing with matching forward & reverse derivatives

so far the triumph of the "mature" tool



- so far the triumph of the "mature" tool
- but then we expanded test cases and ...
- first problem wrong forward values computed:



- so far the triumph of the "mature" tool
- but then we expanded test cases and ...
- first problem wrong forward values computed:
 - originated with partial array initializer list like double a[3]={1.0,2.0};



- so far the triumph of the "mature" tool
- but then we expanded test cases and ...
- first problem wrong forward values computed:
 - originated with partial array initializer list like double a[3]={1.0,2.0};
 - forces initialization in adouble default constructor



- so far the triumph of the "mature" tool
- but then we expanded test cases and ...
- first problem wrong forward values computed:
 - originated with partial array initializer list like double a[3]={1.0,2.0};
 - forces initialization in adouble default constructor
 - can be disabled in Adol-C configuration



- so far the triumph of the "mature" tool
- but then we expanded test cases and ...
- first problem wrong forward values computed:
 - originated with partial array initializer list like double a[3]={1.0,2.0};
 - forces initialization in adouble default constructor
 - can be disabled in Adol-C configuration
 - also changed in Rapsodia



- so far the triumph of the "mature" tool
- but then we expanded test cases and ...
- first problem wrong forward values computed:
 - originated with partial array initializer list like double a[3]={1.0,2.0};
 - forces initialization in adouble default constructor
 - can be disabled in Adol-C configuration
 - also changed in Rapsodia
- second problem wrong forward values computed:



- so far the triumph of the "mature" tool
- but then we expanded test cases and ...
- first problem wrong forward values computed:
 - originated with partial array initializer list like double a[3]={1.0,2.0};
 - forces initialization in adouble default constructor
 - can be disabled in Adol-C configuration
 - also changed in Rapsodia
- second problem wrong forward values computed:
 - originated with more recent location management



- so far the triumph of the "mature" tool
- but then we expanded test cases and ...
- first problem wrong forward values computed:
 - originated with partial array initializer list like double a[3]={1.0,2.0};
 - forces initialization in adouble default constructor
 - can be disabled in Adol-C configuration
 - also changed in Rapsodia
- second problem wrong forward values computed:
 - originated with more recent location management
 - set of fixes (partially by my, partially by KK)



- so far the triumph of the "mature" tool
- but then we expanded test cases and ...
- first problem wrong forward values computed:
 - originated with partial array initializer list like double a[3]={1.0,2.0};
 - forces initialization in adouble default constructor
 - can be disabled in Adol-C configuration
 - also changed in Rapsodia
- second problem wrong forward values computed:
 - originated with more recent location management
 - set of fixes (partially by my, partially by KK)
- overhead?



AdolC-ified ISSM performance - overloading (1)

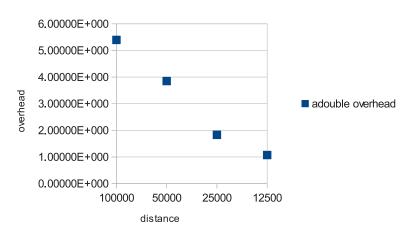
pick some test case (here "test109"), g++ -02, initially horrible timings fixed by another changeset to locations mgmt. from KK



AdolC-ified ISSM performance - overloading (1)

pick some test case (here "test109"), g++ -02, initially horrible timings fixed by another changeset to locations mgmt. from KK

adouble overhead vs max distance





AdolC-ified ISSM performance - overloading (2)

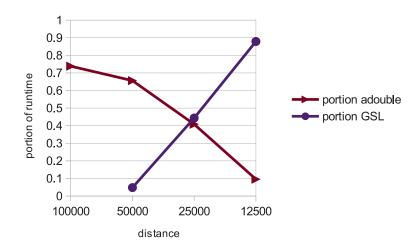
less of a surprise once we look at the portions of runtime



AdolC-ified ISSM performance - overloading (2)

less of a surprise once we look at the portions of runtime

portion of adouble vs GSL over distance





AdolC-ified ISSM performance - tracing & reverse (1)

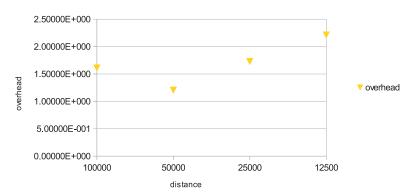
using "test3019" - an AD-enabled regression test



AdolC-ified ISSM performance - tracing & reverse (1)

using "test3019" - an AD-enabled regression test

overhead of tracing and fos_reverse over adouble run





AdolC-ified ISSM performance - tracing & reverse (2)

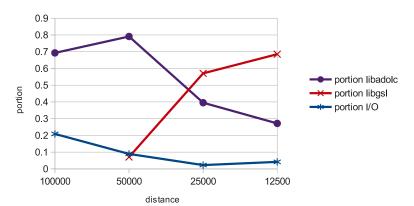
using "test3019" - an AD-enabled regression test



AdolC-ified ISSM performance - tracing & reverse (2)

using "test 3019" - an AD-enabled regression test

portion of total runtime

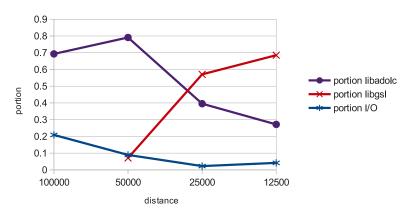




AdolC-ified ISSM performance - tracing & reverse (2)

using "test3019" - an AD-enabled regression test

portion of total runtime



heavily skewed in Adol-C's advantage because of GSL

different solver - MUMPS



- different solver MUMPS
- sparse system



- different solver MUMPS
- sparse system
- parallel setup



- different solver MUMPS
- sparse system
- parallel setup
- had been working on Adjoinable MPI lib (w Laurent & Michel)



- different solver MUMPS
- sparse system
- parallel setup
- ♦ had been working on Adjoinable MPI lib (w Laurent & Michel)
- most of the communication is collective which makes the adjoint "easier"; still needs the AMPI variant called



- different solver MUMPS
- sparse system
- parallel setup
- ♦ had been working on Adjoinable MPI lib (w Laurent & Michel)
- most of the communication is collective which makes the adjoint "easier"; still needs the AMPI variant called
- don't like more preprocessor switches for using MPI or AMPI etc. littering the code



- different solver MUMPS
- sparse system
- parallel setup
- ♦ had been working on Adjoinable MPI lib (w Laurent & Michel)
- most of the communication is collective which makes the adjoint "easier"; still needs the AMPI variant called
- don't like more preprocessor switches for using MPI or AMPI etc. littering the code
- ♦ introduce ISSM_MPI layer to encapsulate 4 versions
 MPI | . | √ | . | √
 AD | . | √ | . |



- different solver MUMPS
- sparse system
- parallel setup
- ♦ had been working on Adjoinable MPI lib (w Laurent & Michel)
- most of the communication is collective which makes the adjoint "easier"; still needs the AMPI variant called
- don't like more preprocessor switches for using MPI or AMPI etc. littering the code
- fake MPI implementation uses memcpy or adouble assignments resp.



- different solver MUMPS
- sparse system
- parallel setup
- ♦ had been working on Adjoinable MPI lib (w Laurent & Michel)
- most of the communication is collective which makes the adjoint "easier"; still needs the AMPI variant called
- don't like more preprocessor switches for using MPI or AMPI etc. littering the code
- fake MPI implementation uses memcpy or adouble assignments resp.
- ♦ layer encapsulates all MPI switching ⇒ cleaner code

- Adol-C has:
 - (i) "tape"-based forward drivers
 - (ii) typed, 4-way stack (tape+Taylors)



- Adol-C has:
 - (i) "tape"-based forward drivers
 - (ii) typed, 4-way stack (tape+Taylors)



- Adol-C has:
 - (i) "tape"-based forward drivers
 - (ii) typed, 4-way stack (tape+Taylors)

makes it "pretty outstanding in terms of uniqueness"

^ Andreas' favorite quote of a Chicago tour guide

(i) means using the same AMPI internal interface as Tapenade



- Adol-C has:
 - (i) "tape"-based forward drivers
 - (ii) typed, 4-way stack (tape+Taylors)

- (i) means using the same AMPI internal interface as Tapenade
- (ii) reluctant to add a 5th stack for MPI parameters with opaque types (such as communicator, datatype etc.) => use Adj.-MPI provided default stack



- Adol-C has:
 - (i) "tape"-based forward drivers
 - (ii) typed, 4-way stack (tape+Taylors)

- (i) means using the same AMPI internal interface as Tapenade
- (ii) reluctant to add a 5th stack for MPI parameters with opaque types (such as communicator, datatype etc.) => use Adj.-MPI provided default stack
- problem with the above is loss of self-containedness of the trace.



- Adol-C has:
 - (i) "tape"-based forward drivers
 - (ii) typed, 4-way stack (tape+Taylors)

- ♦ (i) means using the same AMPI internal interface as Tapenade
- problem with the above is loss of self-containedness of the trace.
- general problems with (re)storing blobs are related to (de)serialization of C++ objects



- Adol-C has:
 - (i) "tape"-based forward drivers
 - (ii) typed, 4-way stack (tape+Taylors)

- (i) means using the same AMPI internal interface as Tapenade
- (ii) reluctant to add a 5th stack for MPI parameters with opaque types (such as communicator, datatype etc.) => use Adj.-MPI provided default stack
- problem with the above is loss of self-containedness of the trace.
- general problems with (re)storing blobs are related to (de)serialization of C++ objects
- covers all interfaces needed by ISSM (reductions, gather/scatter (v) combinations)



- Adol-C has:
 - (i) "tape"-based forward drivers
 - (ii) typed, 4-way stack (tape+Taylors)

- (i) means using the same AMPI internal interface as Tapenade
- (ii) reluctant to add a 5th stack for MPI parameters with opaque types (such as communicator, datatype etc.) => use Adj.-MPI provided default stack
- problem with the above is loss of self-containedness of the trace.
- general problems with (re)storing blobs are related to (de)serialization of C++ objects
- covers all interfaces needed by ISSM (reductions, gather/scatter (v) combinations)
- deals with MPI_INPLACE and 0-count buffers



- Adol-C has:
 - (i) "tape"-based forward drivers
 - (ii) typed, 4-way stack (tape+Taylors)

- (i) means using the same AMPI internal interface as Tapenade
- ◇ (ii) reluctant to add a 5th stack for MPI parameters with opaque types (such as communicator, datatype etc.) ⇒ use Adj.-MPI provided default stack
- problem with the above is loss of self-containedness of the trace.
- general problems with (re)storing blobs are related to (de)serialization of C++ objects
- covers all interfaces needed by ISSM (reductions, gather/scatter (v) combinations)
- deals with MPI_INPLACE and 0-count buffers
- requires contiguous locations ⇒ enforced in xNew spec.

introduction of "active" AMPI_ADOUBLE type to AMPI has been thoroughly discussed between the involved parties



introduction of "active" AMPI_ADOUBLE type to AMPI has been thoroughly discussed between the involved parties

have coexisting active and passive communications



introduction of "active" AMPI_ADOUBLE type to AMPI has been thoroughly discussed between the involved parties

- have coexisting active and passive communications
- buffers passed by void*



- have coexisting active and passive communications
- buffers passed by void*
- forces some external distinction of activity



- have coexisting active and passive communications
- buffers passed by void*
- forces some external distinction of activity
- most natural by corresponding MPI types, particularly when one thinks of derived MPI types



- have coexisting active and passive communications
- buffers passed by void*
- forces some external distinction of activity
- most natural by corresponding MPI types, particularly when one thinks of derived MPI types
- the ISSM wrapper introduces ISSM_MPI_DOUBLE and ISSM_MPI_PDOUBLE



- have coexisting active and passive communications
- buffers passed by void*
- forces some external distinction of activity
- most natural by corresponding MPI types, particularly when one thinks of derived MPI types
- the ISSM wrapper introduces ISSM_MPI_DOUBLE and ISSM_MPI_PDOUBLE
- correspondence required the same way as in standard MPI



- have coexisting active and passive communications
- buffers passed by void*
- forces some external distinction of activity
- most natural by corresponding MPI types, particularly when one thinks of derived MPI types
- the ISSM wrapper introduces ISSM_MPI_DOUBLE and ISSM_MPI_PDOUBLE
- correspondence required the same way as in standard MPI
- practical problems with collectives distributed in the code proved it is easy to get wrong



template the MPI logic with 2 parameters like this pattern

```
#include < iostream >
 234567
      typedef int DataType;
      class TypeInfo {
      public:
        static DataType ourDoubleType;
        static DataType ourIntType;
      };
 8
      DataType TypeInfo::ourDoubleType;
 9
      DataType TypeInfo::ourIntType;
10
11
      template < class T. DataType *typeOfT_p> class C {
12
      public:
13
        C(){};
14
        ~C(){}:
15
        void foo(T aT) { std::cout << aT << "_of_type_" << *typeOfT_p << std::endl; }</pre>
16
      };
17
18
      int main (void) {
19
        TypeInfo::ourDoubleType=1;
20
        TypeInfo::ourIntType=2:
21
        C<double.&TypeInfo::ourDoubleType>().foo(2.0):
22
        C < int, \& TypeInfo::ourIntType > ().foo(-1);
23
        return 0:
24
```

not completed (yet) in ISSM



needs to convey sparsity info to wrapped solver



- needs to convey sparsity info to wrapped solver
- added integer array parameter to a second set of external func interfaces



- needs to convey sparsity info to wrapped solver
- added integer array parameter to a second set of external func interfaces
- suggests void* blobs again but has to address the same concerns as MPI opaque parameters



- needs to convey sparsity info to wrapped solver
- added integer array parameter to a second set of external func interfaces
- suggests void* blobs again but has to address the same concerns as MPI opaque parameters
- question whether to factorize again in the reverse sweep or recover factors:
 - generally tradeoff fill-in for refactoring



- needs to convey sparsity info to wrapped solver
- added integer array parameter to a second set of external func interfaces
- suggests void* blobs again but has to address the same concerns as MPI opaque parameters
- question whether to factorize again in the reverse sweep or recover factors:
 - generally tradeoff fill-in for refactoring
 - specifically MUMPS can dump factors but is written in / geared toward Fortran



- needs to convey sparsity info to wrapped solver
- added integer array parameter to a second set of external func interfaces
- suggests void* blobs again but has to address the same concerns as MPI opaque parameters
- question whether to factorize again in the reverse sweep or recover factors:
 - generally tradeoff fill-in for refactoring
 - specifically MUMPS can dump factors but is written in / geared toward Fortran
 - performance analysis shows it is fast anyway (unlike GSL)

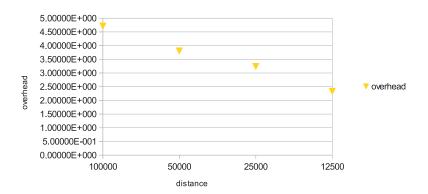


- needs to convey sparsity info to wrapped solver
- added integer array parameter to a second set of external func interfaces
- suggests void* blobs again but has to address the same concerns as MPI opaque parameters
- question whether to factorize again in the reverse sweep or recover factors:
 - generally tradeoff fill-in for refactoring
 - specifically MUMPS can dump factors but is written in / geared toward Fortran
 - performance analysis shows it is fast anyway (unlike GSL)
- ⋄ revisit performance (by now > 500 svn changesets later)



AdolC-ified ISSM performance - tracing & reverse (1)

test3019 - with contiguous locations, 3-way parallel MUMPS overhead of tracing and fos reverse over adouble run

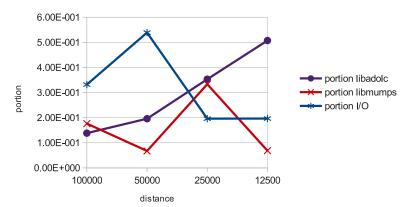




AdolC-ified ISSM performance - tracing & reverse (2)

test3019 - with contiguous locations, 3-way parallel MUMPS

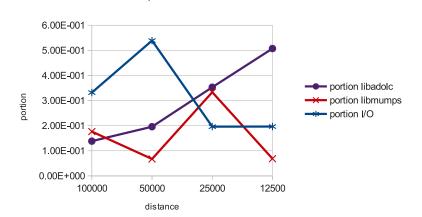
portion of total runtime





AdolC-ified ISSM performance - tracing & reverse (2)

test3019 - with contiguous locations, 3-way parallel MUMPS portion of total runtime

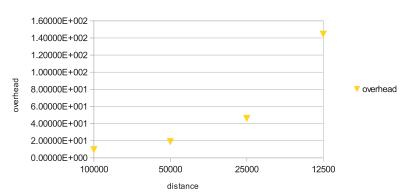


MUMPS is fast - more realistic picture

AdolC-ified ISSM performance - tracing & reverse (3)

test 3019 - with contiguous locations, 3-way parallel MUMPS

ratio total times func+gradient over plain function

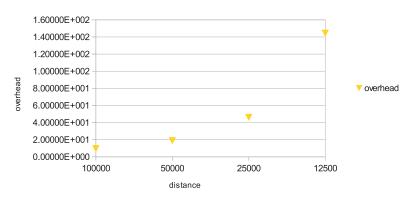




AdolC-ified ISSM performance - tracing & reverse (3)

test3019 - with contiguous locations, 3-way parallel MUMPS

ratio total times func+gradient over plain function



i.e. pretty nasty ... BUT



since these tests happened

identified location management as one of the culprits
 added some logic to constrain the effort in consolidating the locations pool



since these tests happened

- identified location management as one of the culprits
 added some logic to constrain the effort in consolidating the locations pool
- replaced dense (outside of MUMPS) matrix representation with a sparse format



since these tests happened

- identified location management as one of the culprits
 added some logic to constrain the effort in consolidating the locations pool
- replaced dense (outside of MUMPS) matrix representation with a sparse format
- will need a new set of timing tests soon



since these tests happened

- identified location management as one of the culprits
 added some logic to constrain the effort in consolidating the locations pool
- replaced dense (outside of MUMPS) matrix representation with a sparse format
- will need a new set of timing tests soon



since these tests happened

- identified location management as one of the culprits
 added some logic to constrain the effort in consolidating the locations pool
- replaced dense (outside of MUMPS) matrix representation with a sparse format
- will need a new set of timing tests soon

current status:

♦ 60 - way runs on Pleiades



since these tests happened

- identified location management as one of the culprits
 added some logic to constrain the effort in consolidating the locations pool
- replaced dense (outside of MUMPS) matrix representation with a sparse format
- will need a new set of timing tests soon

- ♦ 60 way runs on Pleiades
- transient runs tax the file system and sometimes that causes crashes



since these tests happened

- identified location management as one of the culprits
 added some logic to constrain the effort in consolidating the locations pool
- replaced dense (outside of MUMPS) matrix representation with a sparse format
- will need a new set of timing tests soon

- ♦ 60 way runs on Pleiades
- transient runs tax the file system and sometimes that causes crashes
- acc. to Eric L.: "You have to talk to it with love, and then you get numbers"



since these tests happened

- identified location management as one of the culprits
 added some logic to constrain the effort in consolidating the locations pool
- replaced dense (outside of MUMPS) matrix representation with a sparse format
- will need a new set of timing tests soon

- ♦ 60 way runs on Pleiades
- transient runs tax the file system and sometimes that causes crashes
- acc. to Eric L.: "You have to talk to it with love, and then you get numbers"
- resilience/adjoint integrated checkpointing isn't there yet, but is in the works

since these tests happened

- identified location management as one of the culprits
 ⇒ added some logic to constrain the effort in consolidating
 the locations pool
- replaced dense (outside of MUMPS) matrix representation with a sparse format
- will need a new set of timing tests soon

- ♦ 60 way runs on Pleiades
- transient runs tax the file system and sometimes that causes crashes
- acc. to Eric L.: "You have to talk to it with love, and then you get numbers"
- resilience/adjoint integrated checkpointing isn't there yet, but is in the works
 IOW ... to be continued ...